What is claimed is:

1. A method of addressing a name space mismatch between content servers and content caching systems, comprising steps of:

caching one or more content elements in a content caching system, wherein the cached elements are identified using a first identification format;

automatically learning mappings between a second identification format and the first identification format; and

using the mappings to update one or more selected cached content elements by requesting updated versions of the selected cached content elements from a content server, wherein each request identifies one of the selected cached content elements using the second identification format.

2. The method according to Claim 1, further comprising the step of using the mappings to invalidate one or more particular cached content elements by invoking an invalidation function, wherein each invocation identifies one of the particular cached content elements using the second naming format.

3. A method of addressing a name space mismatch between content servers and content caching systems, comprising steps of:

automatically learning mappings between a first identification format which identifies content elements accessible from one or more content servers and a second identification format which identifies content elements cached by a caching system;

receiving a notification pertaining to one of the cached content elements, wherein the notification identifies the content element using an identifier in the first identification format;

consulting the learned mappings using the identifier in the first identification format from the notification, thereby determining a corresponding identification of the cached content element, wherein the corresponding identification uses the second identification format; and

processing the received notification using the corresponding identification in the second identification format.

4. The method according to Claim 3, wherein:

the identifier in the first identification format identifies more than one of the cached content elements;

the consulting step determines multiple corresponding identifications of the identified cached content elements; and

the processing step processes each of the identified cached content elements using the corresponding identifications in the second identification format.

5. The method according to Claim 3, wherein the notification is a content update notification.

6. The method according to Claim 3, wherein the notification is a content invalidation notification.

7. The method according to Claim 3, wherein the notification is received from a content management system.

8. The method according to Claim 3, wherein a plurality of notifications are received, the plurality comprising at least one content update notification and at least one content invalidation notification, and wherein the plurality of notifications use a consistent interface for conveying the notifications to the content caching system and also to the content servers.

9. The method according to Claim 3, wherein the mappings comprise one or more entries, and wherein the automatically learning step further comprises populating individual ones of the entries with the identification in the first identification format and the corresponding identification in the second identification format in response to receiving the content element to which the entry pertains.

10. The method according to Claim 9, wherein the populating step further comprises steps of:

extracting the identification in the first format from a response message which carries the received content element; and

using this extracted identification, along with the corresponding identification in the second identification format, as values of the populated entry.

11. A method of addressing a name space mismatch between content caching systems and content servers, comprising steps of:

sending a content request from a content caching system to a content server, wherein the content request identifies requested content using a first identification in a first identification format;

receiving a content response from the content server, wherein the content response provides the requested content and is augmented with a second identification of the requested content using a second identification format; and

creating an entry in a mapping, responsive to the receiving step, wherein the entry maps the second identification using the second identification format to the first identification using the first identification format.

12. The method according to Claim 11, wherein the content request signifies an ability to process the augmented content response.

13. The method according to Claim 12, wherein the ability is signified using a message header of the content request

14. The method according to Claim 13, wherein the message header specifies that the caching system is content distribution aware.

15. The method according to Claim 13, wherein the message header is a Hypertext Transfer Protocol ("HTTP") header.

16. The method according to Claim 12, wherein the ability is signified using syntax of a markup language.

17. The method according to Claim 16, wherein the markup language is Hypertext Markup Language ("HTML").

18. The method according to Claim 17, wherein the syntax comprises a "META" tag using an "HTTP-EQUIV" attribute syntax.

19. The method according to Claim 17, wherein the syntax comprises a "META" tag using a "NAME" attribute syntax.

20. The method according to Claim 17, wherein the syntax comprises a specially-denoted comment.

21. The method according to Claim 16, wherein the markup language is XML ("Extensible Markup Language").

22. The method according to Claim 12, wherein the ability is signified using a cookie on the content request.

23. The method according to Claim 11, wherein the augmented content response provides the

second identification using a header.

24. The method according to Claim 23, wherein the header comprises an extended cache control header.

25. The method according to Claim 23, wherein the header comprises a file identification header.

26. The method according to Claim 25, wherein the file identification header is a Hypertext Transfer Protocol ("HTTP") header.

27. The method according to Claim 11, wherein the augmented content response provides the second identification using syntax of a markup language.

28. The method according to Claim 27, wherein the markup language is Hypertext Markup Language ("HTML").

29. The method according to Claim 28, wherein the syntax comprises a "META" tag using an "HTTP-EQUIV" attribute syntax.

30. The method according to Claim 28, wherein the syntax comprises a "META" tag using a "NAME" attribute syntax.

31. The method according to Claim 28, wherein the syntax comprises a specially-denoted comment.

32. The method according to Claim 27, wherein the markup language is XML ("Extensible Markup Language").

33. The method according to Claim 11, wherein the augmented content response provides the second identification using a cookie on the content request.

34. The method according to Claim 11, further comprising the step of caching the received content at the content caching system, responsive to the receiving step.

35. The method according to Claim 34, further comprising the steps of:

receiving an update notification for a selected previously-cached content at the caching system, wherein the update notification identifies the selected previously-cached content using a selected identification in the second identification format;

consulting the mapping, using the selected identification as the second identification in the second identification format, to retrieve the first identification in the first identification format; and

using the retrieved first identification to request an update of the selected previously-cached content.

36. The method according to Claim 35, wherein the using step further comprises steps of:

sending a new content request to the content server, wherein the new content request identifies the requested content using the retrieved first identification;

receiving a new content response from the content server, wherein the new content response provides the requested updated content; and

caching the received updated content.

37. The method according to Claim 34, further comprising the steps of:

receiving an invalidation notification for a selected previously-cached content at the caching system, wherein the invalidation notification identifies the selected previously-cached content using a selected identification in the second identification format;

consulting the mapping, using the selected identification as the second identification in the second identification format, to retrieve the first identification in the first identification format; and

using the retrieved first identification to request an invalidation of the selected previously-cached content.

38. The method according to Claim 11, wherein the first identification format comprises a Uniform Resource Locator ("URL") format.

39. The method according to Claim 11, wherein the second identification format comprises a directory structure format.

40. The method according to Claim 11, wherein the second identification format comprises a file path and file name format.

41. The method according to Claim 11, wherein the second identification is a file path and a file name.

42. The method according to Claim 11, wherein the second identification is a substituted file path and file name.

43. The method according to Claim 42, wherein the substituted file path and file name is an encrypted file path and file name.

44. The method according to Claim 11, further comprising steps of:

receiving, at the content server, the sent content request;

obtaining, by the content server, the requested content;

preparing a response message, by the content server, to transmit the obtained content, wherein the response message is augmented with the second identification in the second identification format; and

sending the prepared response message, including the obtained content, from the content server.

45. A system for addressing a name space mismatch between content servers and content caching systems, comprising:

means for automatically learning mappings between a first identification format which identifies content elements accessible from one or more content servers and a second identification format which identifies content elements cached by a caching system;

means for receiving a notification pertaining to one of the cached content elements, wherein the notification identifies the content element using an identifier in the first identification format;

means for consulting the learned mappings using the identifier in the first identification format from the notification, thereby determining a corresponding identification of the cached content element, wherein the corresponding identification uses the second identification format; and

means for processing the received notification using the corresponding identification in the second identification format.

46. A system for addressing a name space mismatch between content caching systems and content servers, comprising:

means for sending a content request from a content caching system to a content server, wherein the content request identifies requested content using a first identification in a first identification format;

means for receiving a content response from the content server, wherein the content response provides the requested content and is augmented with a second identification of the

requested content using a second identification format; and

means for creating an entry in a mapping, responsive to the means for receiving, wherein the entry maps the second identification using the second identification format to the first identification using the first identification format.

47. A computer program product for addressing a name space mismatch between content servers and content caching systems, the computer program product embodied on one or more computer-usable media and comprising:

computer readable program code means for automatically learning mappings between a first identification format which identifies content elements accessible from one or more content servers and a second identification format which identifies content elements cached by a caching system;

computer readable program code means for receiving a notification pertaining to one of the cached content elements, wherein the notification identifies the content element using an identifier in the first identification format;

computer readable program code means for consulting the learned mappings using the identifier in the first identification format from the notification, thereby determining a corresponding identification of the cached content element, wherein the corresponding identification uses the second identification format; and

computer readable program code means for processing the received notification using the corresponding identification in the second identification format.

48. A computer program product for addressing a name space mismatch between content servers and content caching systems, the computer program product embodied on one or more computer-usable media and comprising, comprising:

computer readable program code means for sending a content request from a content caching system to a content server, wherein the content request identifies requested content using a first identification in a first identification format;

computer readable program code means for receiving a content response from the content server, wherein the content response provides the requested content and is augmented with a second identification of the requested content using a second identification format; and

computer readable program code means for creating an entry in a mapping, responsive to the computer readable program code means for receiving, wherein the entry maps the second identification using the second identification format to the first identification using the first identification format.